

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AIM = 502A

August 1981

CONSTRAINTS

A LANGUAGE FOR EXPRESSING ALMOST-HIERARCHICAL DESCRIPTIONS

BY

GERALD JAY SUSSMAN AND GUY LEWIS STEELE JR.

Abstract:

We present an interactive system organized around networks of constraints rather than the programs which manipulate them. We describe a language of hierarchical constraint networks. We describe one method of deriving useful consequences of a set of constraints which we call propagation. Dependency analysis is used to spot and track down inconsistent subsets of a constraint set. Propagation of constraints is most flexible and useful when coupled with the ability to perform symbolic manipulations on algebraic expressions. Such manipulations are in turn best expressed as alterations or augmentations of the constraint network.

Almost-Hierarchical Constraint Networks can be constructed to represent the multiple viewpoints used by engineers in the synthesis and analysis of electrical networks. These multiple viewpoints are used in terminal equivalence and power arguments to reduce the apparent synergy in a circuit so that it can be attacked algebraically.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. This work was supported in part by the National Science Foundation under Grant MCS77-04828 and in part by Air Force Office of Scientific Research Grant AFOSR-78-3593.

THIS PUBLICATION REPLACES AI MEMO 433 and 502

CONSTRAINTS—A Language for Expressing Almost-Hierarchical Descriptions*

Gerald Jay Sussman** and Guy Lewis Steele Jr.***

Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Mass., U.S.A.

Recommended by Daniel G. Bobrow

ABSTRACT

We present an interactive system organized around networks of constraints rather than the programs which manipulate them. We describe a language of hierarchical constraint networks. We describe one method of deriving useful consequences of a set of constraints which we call propagation. Dependency analysis is used to spot and track down inconsistent subsets of a constraint set. Propagation of constraints is most flexible and useful when coupled with the ability to perform symbolic manipulations on algebraic expressions. Such manipulations are in turn best expressed as alterations or augmentations of the constraint network.

Almost-Hierarchical Constraint Networks can be constructed to represent the multiple viewpoints used by engineers in the synthesis and analysis of electrical networks. These multiple viewpoints are used in terminal equivalence and power arguments to reduce the apparent synergy in a circuit so that it can be attacked algebraically.

1. Introduction

Programming languages are usually organized around unilateral computation. Programs perform predetermined operations on their inputs to produce desired outputs or to evolve a desired process. Physical systems, on the other hand, are usually specified as sets of constraints among several variables. A constraint such as $x*y + z = 3$ is no more about how to compute x given y and z than it is about how to compute z given x and y . Programs are often written to extract useful information from a set of constraints. Such programs depend strongly on the

* This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. This work was supported in part by the National Science Foundation under Grant MCS77-04828, and in part by Air Force Office of Scientific Research Grant AFOSR-78-3593.

** Jolly Good Fellow

*** Fannie and John Hertz Fellow

forms of the constraints and are committed to solving for particular variables. For example, programs for electrical circuit analysis are usually not good for circuit synthesis as well.

We present an interactive system organized around networks of constraints rather than the programs which manipulate them. We describe a language of hierarchical constraint networks. It supplies a set of primitive constraint types and a means of building compound constraint systems by combining instances of simpler ones. Constraint networks can be constructed which violate hierarchy in controlled ways. This is useful for representation of the multiple (perhaps redundant) views of a system which characterize much of our understanding of engineered systems.

The constraint language was developed for representing knowledge about electrical circuits.¹ The ability to represent multiple redundant views is critical for implementing SLICES, a device we invented for notating some of the strategies that we observed are used by expert circuit designers. In circuit synthesis one problem is to determine the component values in a circuit of known topology and given specifications. This is in general a very difficult algebra problem. It is infeasible to determine the component values by solving the set of equations and inequalities which result from matching the symbolic analysis of a circuit with the given design parameters. An expert circuit designer uses terminal equivalence and power arguments to reduce the apparent synergy in a circuit and focus his attention. At any instant an expert will only concentrate on a small portion of the circuit, assuming that the rest will "work as planned". This assumption is captured by summarizing the behavior of the part of the circuit peripheral to the area under attention as a set of equivalences. Equivalences are also used to summarize the specifications of the behavior of the part under attention. This allows the designer to reason about the details of one section of the circuit in isolation. Few assumptions about the nature of adjacent parts of a circuit are allowed to migrate across the boundary of equivalences. From the point of view of circuit analysis SLICES model the way an experienced designer uses terminal equivalences to provide him with multiple local views of a circuit.

We also describe one method of deriving useful consequences of a set of constraints which we call propagation.² Propagation automatically takes advantage

¹ This work is part of a larger effort to apply artificial intelligence methods to computer-aided design at the MIT AI Laboratory. We are not solely interested in computer-aided circuit design [35, 36], but rather in an understanding of the general epistemology of engineering. We are also engaged in a substantial effort on the computer-aided design of software systems [27] and computer-aided design of VLSI chips [37].

² 'Propagation of constraints' was originally invented as a generalization of 'Guillemin's method' of analyzing electrical ladder circuits. It was used in the analysis programs EL [40] and ARS [38], and in the synthesis program SYN [8]. The basic idea of the method was first described in [2] as part of a method for localizing faults in electrical circuits. De Kleer also used propagation analysis in his fault localizer [7]. Sutherland [42] appears to have developed a similar technique (the "One Pass Method") for constraint satisfaction in Sketchpad.

of the sparseness of most constraint networks. Propagation is not prematurely committed to solving for any particular predetermined set of unknowns, nor is it committed to using any particular constraint in a predetermined way. In fact propagation analysis may be freely intermixed with the addition of new constraints and the deletion of old ones. This requires that the system keep track of how its conclusions (and intermediate results) are consequences of the particular constraints they were derived from. Dependency analysis is also used to spot and track down inconsistent subsets of a constraint set. Propagation of constraints is most flexible and useful when coupled with the ability to perform symbolic manipulations on algebraic expressions. Such manipulations can be expressed as alterations or augmentations of the constraint network.

2. The Language of Constraints

A language is a means of communication of ideas. A language generally has a "theme," the class of ideas which it is optimized for communicating. For example, most computer languages are designed for expressing algorithms. They are optimized for communicating imperative, procedural notions.³ The theme of the constraint language is declarative. It is good for expressing descriptions of structural relationships and physical situations.


Every language has a set of primitive entities which represent the elementary notions of the domain, and rules of composition by which compound entities are constructed.

Simple Constraints

We will now proceed to illustrate the entities of the constraint language by an annotated interaction with the constraint-language interpreter. Input to the interpreter is in lower case, following the prompt '==> .'

One kind of primitive (i.e., built-in) constraint we might have is an 'adder' which constrains three numbers (called the addend, augend, and sum of the adder) in such a way that the addend plus the augend must equal the sum. Put another way, it constrains the augend to equal the difference of the sum and the addend. (The point is that there is no preferred direction of computation.) First, let's make an adder, called ZAP:


==> (create zap adder)
ZAP



³ Some examples of languages which are less organized around the notion of procedure are such simulation systems as DYNAMO, GPSS, and SIMULA.

Every object in the constraint language has a type, and may have parts. We can ask ZAP to tell us about itself:

```
==> (>> type? zap)
ADDER
==> (>> partnames? zap)
(A1 A2 SUM)
```



It seems that ZAP has three parts, the addend (A1), the augend (A2), and the sum (SUM). (A form such as (>> a b c) may be read as 'the a of the b of c.' Thus it constitutes a 'pathname' for selecting an item starting at the root c.) The parts also have types:

```
==> (>> type? a1 zap)
CELL
==> (>> partnames? a1 zap)
NIL
```


A cell is a primitive entity which hath no parts.⁴ Cells are used for two things in the constraint language. They are used to hold computational values, and (as we will see later) they can be connected together when building compound constraints. We can use WHAT-IS to discover the value in a cell:

```
==> (what-is (>> sum zap))
Sorry, I don't know it. I need:
(>> A1 ZAP)
(>> A2 ZAP)
to use rule: (>> ADDER-RULE#1 ZAP)
```

At this point, the system does not know a value for the SUM of ZAP. However, it has told us that if we had given it the A1 and A2 it could have computed a sum using a rule (ADDER-RULE#1) internal to ZAP. Instead, we now give the sum a value.

```
==> (set-parameter (>> sum zap) 5.0)
5.0
```

```
==> (what-is (>> sum zap))
(>> SUM ZAP) = 5.0
==> (why (>> sum zap))
Because you told me so.
```




⁴ Euclid had a point here.

Now that the system knows about the SUM of ZAP, can it tell us about the A1?

```
==> (what-is (>> a1 zap))
Sorry, I don't know it. I need:
(>> A2 ZAP)
to use rule: (>> ADDER-RULE#2 ZAP)
```

The A1 is not yet completely determined, but it could be determined if the A2 were known. If instead we specify the A1, the A2 should be determined.

```
==> (set-parameter (>> a1 zap) -2.0)
-2.0
==> (what-is (>> a2 zap))
(>> A2 ZAP) = 7.0
==> (why (>> a2 zap))
I used rule (>> ADDER-RULE#3 ZAP) on:
(>> SUM ZAP)
(>> A1 ZAP)
```




We have just illustrated the use of an adder, a typical entity of the constraint language. A constraint enforces a relationship among several entities. If enough information is known to immediately deduce unknown cell values, those values are computed. These new values may enable further deductions. We call this deductive process 'propagation of constraints.' (Actually, it is values that are propagated, through a network of constraints.)

Networks of constraints

Suppose we also have a kind of constraint called a 'multiplier':

```
==> (create foo multiplier)
FOO
```



Neither an adder by itself nor a multiplier by itself is a particularly fascinating device. The system provides a means for combining simple constraints into arbitrarily complicated networks.

In typical algorithmic computer languages a compound algorithm is built from simpler ones by rules of temporal sequencing and data flow. Algorithms can be temporally concatenated (sequencing), selected among (conditionals), and iterated (do loops). Data flow is indicated explicitly by functional composition and implicitly by side effects on shared variables.

In the constraint language, a compound constraint is built from simpler ones by linking some of their parts. The method of connection is a declaration that two entities are in fact the same.

The '=' primitive allows us to identify any two constraint entities of the same type. In particular, if two cells are identified they are effectively considered to be

the same cell—they are constrained to have the same value.

```
==> (== (>> product foo) (>> a1 zap))
```

IDENTITY

```
==> (== (>> m1 foo) (>> a2 zap))
```

IDENTITY

We have linked various parts of our adder ZAP and our multiplier FOO. The resolving network looks like Fig. 1.




FIG. 1.

Now ZAP had some values in its cells when we linked them to FOO's cells. These values have already been propagated through FOO.

```
==> (what-is (>> m2 foo))
```

```
(>> M2 FOO) = -0.285714287
```

```
==> (why m2 foo))
```

I used rule (>> MULTIPLIER-RULE #5 FOO) on:

```
(>> PRODUCT FOO)
```

```
(>> M1 FOO)
```

We can chase these deductions to their ultimate reasons.

```
==> (why (>> product foo))
```

I used rule (>> 1 <= 2) on:

```
(>> A1 ZAP)
```

The rule '1 <= 2' is a manifestation of the linkage between the PRODUCT of FOO and the A1 of ZAP.

```
==> (why (>> a1 zap))
```

Because you told me so.

```
==> (why (>> m1 foo))
```

I used rule (>> 1 <= 2) on:

```
(>> A2 ZAP)
```

```
==> (why (>> a2 zap))
```

I used rule (>> ADDER-RULE #3 ZAP) on:

```
(>> SUM ZAP)
```

```
(>> A1 ZAP)
```

It is often convenient to be able to determine the ultimate antecedents of a deduced

value. For example, we know the value of the M2 of FOO because of a chain of deductions ultimately derived from our knowledge of the SUM and A1 of ZAP.

```
==> (premises (>> m2 foo))
(>> SUM ZAP) = 5.0
(>> A1 ZAP) = -2.0
```

If we change a premise all conclusions which depended upon it are automatically retracted and new conclusions are drawn. In a large network only a few values may depend on any one premise. The constraint language interpreter retracts only those values and only the incremental new deductions are made (see Fig. 2).

```
==> (change-parameter (>> a1 zap) 2.0)
2.0
```

```
==> (what-is (>> a2 zap))
(>> A2 ZAP) = 3.0
```

```
==> (what-is (>> m2 foo))
(>> MS FOO) = 0.666666664
```




FIG. 2.

We cannot so easily change a parameter whose value is a consequence of other known facts. One of the premises upon which the changing parameter depends must be abandoned. The system automatically chases down only the relevant premises and gives us a choice of which we want to retract.

```
==> (change-parameter (>> a2 zap) 19.0)
Which of the following assumptions will you change?
1 (>> A1 ZAP) = 2.0
2 (>> SUM ZAP) = 5.0
ANSWER:
```

Before allowing the system to change anything, we may want to investigate the repercussions of changing it. (Note that the prompt has changed to '-->', indicating that the system is waiting for an answer to its question.)

```
--> (results (>> sum zap))
Rule (>> ADDER-RULE #3 ZAP) got (>> A2 ZAP) = 3.0

--> (results (>> a1 zap))
Rule (>> 1 <= 2) got (>> PRODUCT FOO) = 2.0
RULE (>> ADDER-RULE #3 ZAP) got (>> A2 ZAP) = 3.0
```

We decide to retract assumption 1, that the A1 of ZAP is 2.0.

--> (answer 1)
19.0

The system finished changing the parameter to 19 and restored the prompt to '==>'. The new value of the M2 of FOO is now:

==> (what-is (>> m2 foo))
(>> MS FOO) = -0.7368421

The new situation is as shown in Fig. 3.




FIG. 3.

Local constraint propagation cannot solve every problem of assigning values to cells in our network. Sometimes a more global analysis is necessary. For example, if we change the value of the M2 of FOO to 1.0 and we release the premise concerning the value of the A2 of ZAP, the only consistent value that either the A1 or the A2 of ZAP can take on is 2.5. (See Fig. 4.)

==> (change-parameter (>> m2 foo) 1.0)
Which of the following assumptions will you change?
1 (>> A2 ZAP) = 19.0
2 (>> SUM ZAP) = 5.0

ANSWER:

--> (answer 1)
1.0

==> (what-is (>> a1 zap))
Sorry, I don't know it. I need:

(>> PRODUCT FOO)

to use rule: (>> 2 <= 1)

or, I need:

(>> A2 ZAP)

to use rule: (>> ADDER-RULE#2 ZAP)




FIG. 4.

The problem is that neither the adder nor the M2 has enough information to make any deductions by itself. However, if they could collaborate, they could come to a correct conclusion. We will discuss this later.

Inasmuch as the constraint network cannot determine the PRODUCT of FOO, we will specify a value explicitly.

==> (change-parameter (>> product foo) 17.0)

Contradiction . . . disputants are:

(>> PRODUCT FOO) = 17.0

(>> CHALLENGER: MULTIPLIER-RULE #3 FOO) = -12.0

Choose an assumption to change:

1 (>> M2 FOO) = 1.0

2 (>> SUM ZAP) = 5.0

ANSWER:




FIG. 5.

Unfortunately, the very act of making the assumption that the PRODUCT of FOO is 17.0 allows the deduction that the PRODUCT of FOO must (also) be -12.0, given the other assumptions in force. This immediate contradiction is noticed, and the system presents us with a choice of actions.

--> (answer 2)

17.0

We have chosen to discard assumption 2, and the network readjusts itself. (See Fig. 6.)

==> (what-is (>> sum zap))

(>> SUM ZAP) = 34.0

==> (what-is (>> a2 zap))

(>> A2 ZAP) = 17.0




FIG. 6.

Now we have seen the rule of composition by which constraints are combined into networks of constraints by identification of parts. We have also seen how it is possible to compute with constraints. The constraint interpreter maintains the 'dependencies' which describe how conclusions are drawn from their antecedents. Besides the fact that the dependencies are important to the system for efficient computation (as we shall see), they are useful for debugging purposes, and for helping a user understand the behavior of his constraint network. We believe that this is a good idea. A discipline should be developed for construction of systems so that they are responsive to their users. They must be *responsible for their conclusions* and able to explain how the conclusions were derived. This is increasingly important as systems become more complex, thus exceeding the understanding of any one person.

Abstraction and hierarchy

Using the '==' construction, we can build networks of adders and multipliers which are arbitrarily large, and thus arbitrarily complex. In order to deal with this complexity, we need a way to break up large networks into meaningful pieces. If we are lucky, many of these pieces will be the same, and can be considered single building blocks at a higher conceptual level. We need a way to define such building blocks in terms of smaller ones.

We would like these compound building blocks to behave in the same way as the primitive entities of the language. This allows us to combine such building blocks using the same construction methods. In this way we can build arbitrarily complicated compound objects in a hierarchical manner. The hierarchy allows the complexity at any one level to be limited.

In languages which permit this kind of hierarchical definition there is a rule of abstraction, which specifies how a compound object is to be treated as a simple one. Such a rule must specify a particular combination of simpler objects which is to be abstracted, and an interface between the specified combination and the external appearance of the abstraction as a single object.⁵

The constraint language provides a rule of abstraction which allows us to associate a name with a pattern designating component constraints and '==' specifications for linking them. When a pattern is instantiated, the designated components are recursively instantiated and then linked according to the '==' specifications.

For example, we can combine adders and multipliers to produce a kind of compound constraint called a 'resistor.' An (ideal) resistor is essentially an object which enforces numerical constraints among two node potentials, two currents,

⁵ A simple example of such a rule is that of lambda-abstraction [4], in which a lambda-expression packages up a possibly compound expression and an interface specification in the form of a variable name (in LISP, variable names).

and a resistance:

$$(v_1 - v_2) = i_1 R$$

$$i_1 + i_2 = 0$$





We notate this as follows:

==> (constraint resistor

```

((v1 number)
 (v2 number)
 (i1 number)
 (i2 number)
 (resistance number)
 (av adder)
 (ai adder)
 (m multiplier))
(== v1 (>> sum av))
(== v2 (>> a1 av))
(== (>> a2 av) (>> product m))
(== i1 (>> m1 m))
(== resistance (>> m2 m))
(== i1 (>> a1 ai))
(== i2 (>> a2 ai))
(constant (>> sum ai) 0.0))

```



RESISTOR

The keyword 'constraint' is followed by a name, a list of component names and types, and a set of linkages. The expression (constant <something> <value>) is an abbreviation for forcing <something> to have the given computational value.⁶


⁶ Our method of specifying an object by a set of subparts and a set of identifications of sub-subparts is essentially the same as that independently developed by Alan Borning, a graduate student at Stanford University. In his Ph.D. thesis, Stanford University (July 1979), entitled "ThingLab—A Simulation Laboratory," Borning develops an interactive system (written in Smalltalk [16]) for simulating the effects of constraints. ThingLab provides a beautiful and convenient graphics interface controlled by the constraint network. In contrast to the system described here, ThingLab does not retain dependency information, and uses relaxation techniques (rather than algebra or multiple redundant views) to deal with systems not directly amenable to simple propagation. In these respects ThingLab is very similar to Sketchpad [42]. Unlike Sketchpad, ThingLab provides facilities for incrementally compiling constraints, as well as a non-graphical, programming-language notation for the constraints. [1] is a preliminary description of ThingLab's capabilities. What we have referred to as identification of parts Borning, following Sutherland, calls 'merging.' This is because in both systems data structures representing identified objects are actually merged internally to become a single data structure. This is not done in our constraint system because the fact of such an identification must be explicitly recorded for the sake of the dependency information (which is not included in the other two systems).

With this definition we can make a resistor in the same way that we make an adder:

```
==> (create r43 resistor)
R43
```


This causes R43 to be the name of an *instance* of the pattern RESISTOR. In instantiating the pattern, instances of components such as adders and multipliers are also effectively created. We can use instances of resistor to build a voltage divider:

```
==> (create r44 resistor)
R44
==> (create kcl1 adder)
KCL1
==> (== (>> v2 r43) (>> v1 r44))
IDENTITY
==> (== (>> i2 r43) (>> a1 kcl1))
IDENTITY
==> (== (>> i1 r44) (>> a2 kcl1))
IDENTITY
==> (constant (>> sum kcl1) 0.0)
0.0
```



It is silly to use adders and numerical connectors to hook together resistors. We should not be mixing conceptual levels. In the world of electrical circuits, elements like resistors have *terminals* which connect to *nodes*. A terminal has a potential ('voltage') on it and a current into it, packaged up together. A node connects two or more terminals so as to constrain their potentials to be the same, and the sum of the currents into them to be zero.


```
==> (constraint terminal ((v number) (i number)))
TERMINAL
```



The 'terminal' constraint is trivial. All it does is package up two numbers so that they can be referred to as a unit.

```
== > (constraint 2-node
      ((t1 terminal)
       (t2 terminal)
       (kcl adder))
      (== (>> v t1) (>> v t2))
      (== (>> i t1) (>> a1 kcl))
      (== (>> i t2) (>> a2 kcl))
      (constant (>> sum kcl) 0.0))
```

2-NODE



A '2-node' is a two-terminal node. It has an adder for enforcing the sum-of-currents constraint (Kirchoff's Current Law). The equal-potentials constraint (Kirchoff's Voltage Law) is enforced by direct identification.

==> (constraint resistor

((t1 terminal)

(t2 terminal)

(resistance number)

(av adder)

(ai adder)

(m multiplier))

((= (>> v t1) (>> sum av))

((= (>> v t2) (>> a1 av))

((= (>> a2 av) (>> product m))



((= (>> i t1) (>> m1 m))

((= resistance (>> m2 m))

((= (>> i t1) (>> a1 ai))

((= (>> i t2) (>> a2 ai))

(constant (>> sum ai) 0.0))



This definition is as before, except that terminals are used.

Now we can simply create two resistors and a 2-node, and connect them up by identifying the terminals of resistors with terminals of the node.

==> (create r43 resistor)

R43

==> (create r44 resistor)

R44

==> (create n2 2-node)


N2

==> ((= (>> t1 r44) (>> t2 n2))

IDENTITY

==> ((= (>> t2 r43) (>> t1 n2))

IDENTITY



In the constraint language, it is possible to identify compound constraints, using '=' in exactly the same way that primitive objects like numbers are identified. The meaning of equality for compound objects is that two objects are equal if their corresponding components are equal. Thus when two terminals are made equal, their voltages are identified, and so are their currents.

(In the examples involving electrical components we have been, and will be, carefully skirting a difficult issue regarding the directions of currents. In the constraint language, things 'connected together' are identified as being the same thing. In the electrical world, while electrically connecting two electrical terminals makes their voltages the same (Kirchoff's Voltage Law), it does not make their

currents the same. The currents are instead subject to the more complex constraint that they must sum to zero. This is the reason we use nodes as an intermediate connector for enforcing Kirchhoff's Current Law. Nodes must be used carefully to keep the current directions straight. Two constraints describing electrical devices cannot be directly connected, but must be connected through a node. Additionally, two nodes cannot be directly connected, except through a constraint representing a device. The examples we give here are correct, but it is very easy to misuse the language.)

We can use this structure to compute parameters of the circuit. The two resistors form a voltage divider. We can, for example, specify the voltages at the ends and middle of the divider, and the resistance of one resistor. The algebraic constraints of the circuit will then permit the deduction of the resistance of the other resistor:

```
==> (set-parameter (>> v t2 r44) 0.0)
0.0
==> (set-parameter (>> v t1 r44) 3.0)
3.0
==> (set-parameter (>> v t1 r43) 10.0)
10.0
==> (set-parameter (>> resistance r44) 9.0)
9.0
==> (what-is (>> resistance r43))
(>> RESISTANCE R43) = 21.0
```

This is similar to computations we have already seen. The 'real work' is being done by adders and multipliers which are part of the description we have provided of resistors.

As before, there can be problems with not being able to compute something with constraints for which we in principle have enough information. For example, given the voltage at the top and the bottom, and the resistances of both resistors, it should be possible to compute the voltage at the midpoint (the 'divided voltage'):

```
==> (forget-parameter (>> v t1 r44))
NIL
==> (set-parameter (>> resistance r43) 21.0)
21.0
==> (what-is (>> v t1 r44))
Sorry, I don't know it. I need:
(>> V T2 N2)
to use rule: (>> 1 <= 2)
or, I need:
(>> SUM AV R44)
to use rule: (>> 1 <= 2 R44)
```

Forgetting the voltage at the midpoint caused the system also to forget the resistance

of R43, which had been deduced from this voltage. Reinstating the resistance value should permit the deduction of the midpoint voltage, but as before this cannot be done by the constraints. There are several ways around this problem. One is to introduce multiple views of the circuit; we will examine this technique in the next section. Another is to use symbolic algebra in the computations; this will be examined later.

3. Almost-Hierarchical Systems

To be powerful, a description must impose a conceptual structure on a system for the purpose of limiting the complexity in understanding or designing it. If a system cannot be immediately understood in its entirety (by 'Gestalt') then it must be understood in pieces. The descriptive structure directs attention to sets of components of the system which together constitute a single conceptual piece.

Often a system can be partitioned into pieces which are more or less disjoint and which together cover the entire system. The total system can be understood by understanding the pieces and by understanding the composition by which the pieces constitute the system. Similarly, each piece may be similarly partitioned. In this way we derive a single tree-like decomposition of the system. Such a tree we call a hierarchy.

It is likely, however, that at any stage there is more than one useful partitioning of a piece. If so, then a single hierarchy does not suffice to indicate all the conceptual pieces of interest in the system. Pieces whose sub-pieces are localized within one hierarchy will probably have its sub-pieces widely dispersed throughout another.

For example, a mechanical timepiece has among its major parts a mainspring, several wheels or 'gears' (main, center, third, fourth, hour, minute, and escape wheels), a lever, a balance/hairspring assembly, and hands. (See Fig. 7.) These parts are usually grouped into sets which are physically connected: the mainspring and main wheel together form the barrel; the center, hour, and minute wheels plus the hands form the motion works; the main, center, third, fourth, and escape wheels form the wheel train; and so on. Also, some of the parts are themselves




FIG. 7. Structural decomposition of a watch.




FIG. 8. Functional decomposition of a watch.

complicated assemblies, for example the balance. Notice that there is some overlap between the main structural groupings where they interface.

A timepiece also may be *functionally* described as a fivefold decomposition, as in Fig. 8. Conceptually these five functional parts are disjoint. In every real mechanical timepiece, however, the same physical part may serve in more than one functional unit. The wheel train which transmits energy from the mainspring to the escapement also performs part of the scaling between the escapement's oscillations and the display (dial); but the wheel train is *not* part of the oscillator module which appears between the transmission and scaling modules. The escape wheel is part of both the oscillator module and the wheel train. Other parts are also shared among modules.

In Fig. 9, the structural and functional hierarchies are (partially) shown together. Notice that each is not a strict hierarchy, but almost is; and that the two almost-hierarchies are not the same, but are similar. In some places a single unit in one hierarchy is the same as a single unit in another (for example, the wheel train and transmission coincide). In other places what is a single unit in one hierarchy is spread out in the other (for example, the single functional idea of scaling is spread out in the structural hierarchy, while the wheel train and motion works structural units each have various purposes in the functional hierarchy).

It would be possible to design a timepiece which exhibited much less structural overlap between functional units, so that the structural and functional hierarchies would be almost identical. While such a timepiece would certainly exhibit 'structured' design, in the sense of 'structured programming,' it would be much less

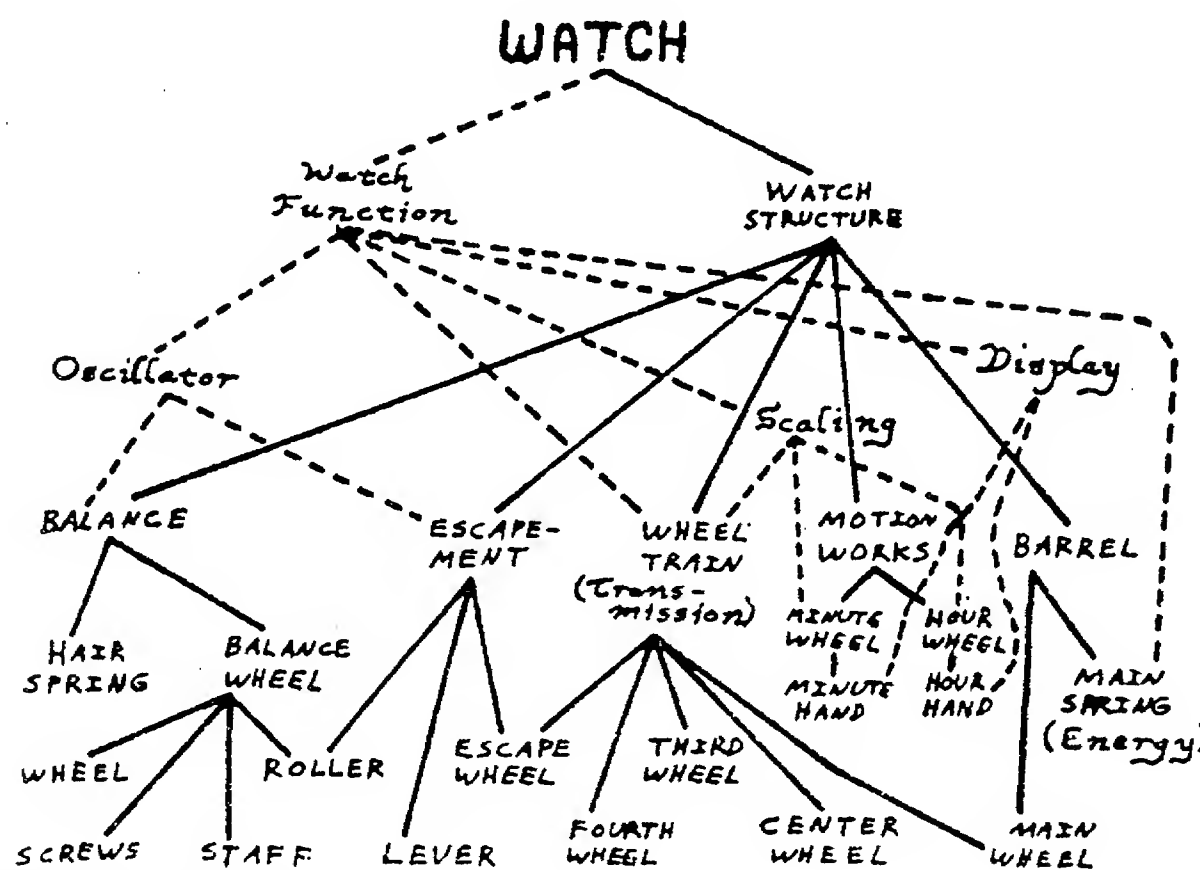


FIG. 9. Overlapping decompositions of a watch.

reliable (due to unnecessary duplication of parts to avoid overlap), and would be very difficult to make small enough to ride comfortably on one's wrist.

Other engineered systems, such as electrical circuits and computer programs, exhibit a similar sharing of parts among uses. A strictly hierarchical description can be only an approximation to the true structure of an object. Conversely, when designing an object, a good engineer may start with a hierarchical plan, but will then slightly destroy the hierarchy when interfacing the substructures. This occurs because any particular hierarchical description emphasizes some groupings of features at the expense of other equally valid or important ones. A given hierarchy may be violated because of constraints not visible within that hierarchy.

A good description explicitly acknowledges the sharing of structure in engineered devices. This can be captured as a locally hierarchical structure in which each feature is described and related to 'lower-level' features which implement that feature, even though those sub-features may be spread throughout some other (also locally hierarchical) part of the description and used to implement other features as well.

The constraint language naturally permits one to express almost-hierarchical descriptions, because its rules of composition are explicitly stated in terms of the sharing of parts. The rule of abstraction is hierarchical, but the rule of composition permits one to violate the hierarchy in a controlled manner.

In fact, we have already used this feature to describe the connection of two resistors. One terminal of each resistor is shared with a terminal of the node. Terminals are not very interesting as they have no non-trivial internal structure. However, let us extend that example to formalize a more interesting almost-hierarchy.

Equivalence slices

Electrical Engineering has provided us with the key concept of an equivalent circuit and several theorems which tell us about some useful types of equivalence⁷ but it does not supply us with a notation for drawing the simultaneous multiple views of a circuit that make the use of equivalence so powerful. The best that one ever sees in engineering books is block diagrams, but the boxes we wish to draw would be so overlapped as to be difficult to draw in that way. We now introduce SLICES, a notation for formally describing these simultaneous views, so that they can be communicated to either a person or a computer involved in the design or analysis of circuits.⁸

⁷ In electrical engineering, most equivalences, such as the Thevenin/Norton theorems and the two-port results depend critically on the linearity of the circuits they summarize. Thus they are severely limited in applicability. Engineers often compensate for this limitation by considering perturbations. This 'small signal analysis' is critical to the design of complex signal processing systems.

⁸ One difficulty with current teaching in Electrical Engineering (and other disciplines) is the lack of formal ways of communicating 'intuitive' knowledge. Much of this knowledge is procedural in nature. The spread of 'computer culture' will enhance our ability to express this kind of knowledge. Slices are a way of expressing more traditional intuitive knowledge about how a circuit works.

Usually in thinking of electrical networks we only consider one kind of connection—Terminals of components are ‘soldered’ together to form nodes which have definite node voltages and which conserve current. If we want to specify, in a schematic diagram, both a circuit and its equivalent, we cannot just connect them with the usual kind of solder because they would indicate that they are connected in parallel. We are forced to introduce a new kind of connection: Let us define 2 ‘wires’ as being *identified* if they have the same voltage and the *same* current (in corresponding directions). Thus, we can indicate that two circuits are equivalent by identifying the corresponding terminals. For example, Fig. 10 indicates that network N has a Thevenin equivalent.

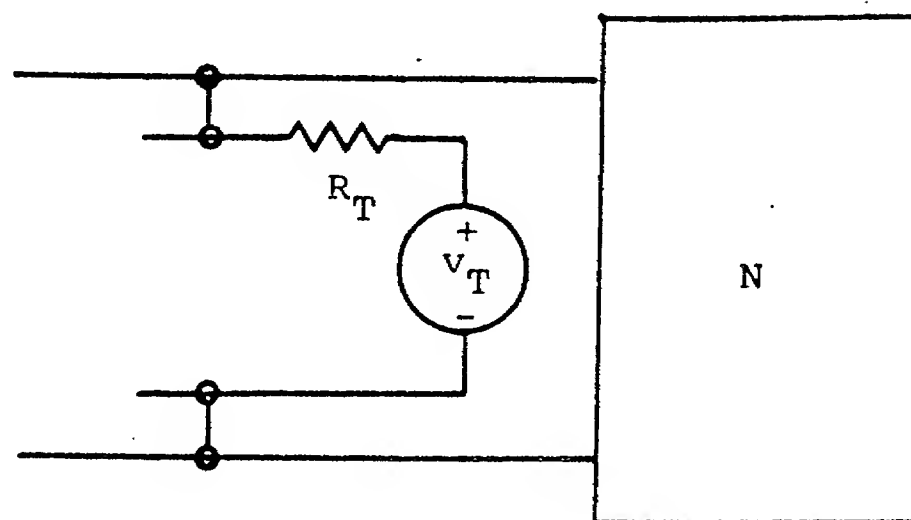


FIG. 10. Thevenin equivalent.

Any particular circuit diagram may contain several slices.⁹ Each slice involves two or more identifications of wires. Thus, to keep things straight, we label each identification with the name of the slice it is part of. Thus network N may have both a Thevenin (slice a) and Norton (slice b) equivalent. (See Fig. 11.)

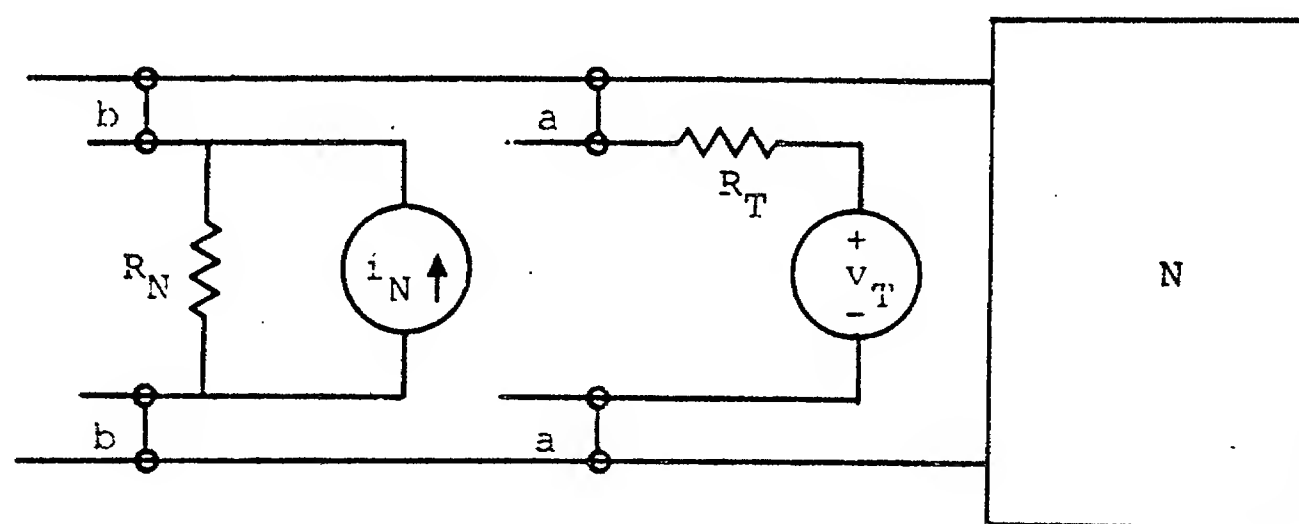


FIG. 11. Thevenin and Norton equivalents.

⁹ Stallman and Sussman [38] introduced the ‘grey box,’ a predecessor of the slice idea. Grey boxes also attached to a circuit diagram by identification of terminals. Grey boxes are additional ‘laws’ which bypass the deductive machinery and present the answer for certain otherwise unfeasible deductions. For example, the incremental gain of an emitter-coupled pair is impossible to deduce without exponential algebra. These additional laws allow EL to ‘know’ some facts about a circuit construct in an implicit manner. Grey boxes, however, have no substructure so they cannot really provide an alternate electrical view of a circuit. Slices, on the other hand are ‘real’ electrical subdiagrams. They can contain further slices describing their structure. In addition, slices can share parts and can identify wires in two different other slices.

Parameter slices

Identifications of wires is not enough to express many constraints. Sometimes we need to be able to say that two circuit parameters are to be identified. For example, if two resistances are constrained to be the same or if the power dissipated in one part is the same as the power dissipated in some other part. We use dotted lines to express identifications between parameters. For example, to express the truth that $R_T = R_N$ in our network above we would draw Fig. 12.




FIG. 12. Shared parameters.

Here we have an identification which spans two slices. This is perfectly legitimate. Once we have placed a slice on the circuit diagram it becomes part of the diagram and further slices can include parts from it. We cannot construct, however, two distinct slices, each of which refers to parts 'inside' the other.

The problem of nodes

There is another problem in standard electrical circuit notation which has to be overcome. In formal network theory the diagrams of Figs. 13(a) and 13(b) are considered identical. This identity is not reasonable to one who is interested in the current, i , indicated on diagram a. The topological abstraction of a node which conserves current and which has a definite potential is not the correct representation of the circuit from his point of view. We thus consider *connection points* to be components with simple laws, and we give *wires* formal status. Thus we may have several equivalent circuits of the 'same' (network theoretic) node which are used when we wish to talk about different groupings of the currents.




FIG. 13.


Slices as constraint diagrams

We can express the idea that two resistors in series are equivalent to a single resistor:

```

==> (constraint series-resistors
      ((r1 resistor)
       (r2 resistor)
       (req resistor)
       (a adder))
      (== (>> t1 r1) (>> t1 req))
      (== (>> t2 r2) (>> t2 req))
      (== (>> a1 a) (>> resistance r1))
      (== (>> a2 a) (>> resistance r2))
      (== (>> sum a) (>> resistance req)))

```

**SERIES-RESISTORS**

(Remember: This does not say that (>> t1 req) is electrically connected to (>> t1 r1). It says that (>> t1 req) *is* (>> t1 r1). This differs from the situation where two resistors are connected through a node. Here we are talking about two views of the same terminal.)

This definition does not express the fact that we ordinarily apply it only to resistors which share a common node. This is a restriction on *usage* of the constraint, rather than a restriction to be enforced *by* the constraint. If it is used properly, however, it will enforce the desired series constraint. We now construct an instance of the constraint, and use it to describe the resistors R44 and R43 connected earlier.

```

==> (create srl series-resistors)
SR1
==> (== (>> r1 srl) r43)
IDENTITY
==> (== (>> r2 srl) r44)
IDENTITY

```

Now there are two simultaneous, redundant descriptions of what is going on between the upper and lower terminals of the circuit. Neither description is regarded as being 'more real' than the other; we merely say that one or the other is more useful for some purpose. Both descriptions, as one resistor and as two, are equally valid. We can speak of (>> resistance req srl), the resistance of the resistor in the single-resistor description, just as well as of (>> resistance r43) or (>> i t1 r44), which are quantities of the two resistor description.

These redundant descriptions have important computational applications. For example, in the voltage divider circuit the alternative description of the divider as a single resistor permits the deduction of the current through the divider. This in turn permits the deduction of the voltage at the midpoint!


```

==> (what-is (>> v t1 r44))
(>> V T1 R44) = 3.0

```

What we have done here is to introduce an *alternative point of view* of the circuit. While in principle it contains no extra information, the new viewpoint is better organized for certain purposes (such as determining the current through the divider). The constraint language permits us to introduce many such redundant viewpoints so that they can cooperate in solving a problem.

Let us add another resistor R45 to our circuit:



We can describe various series combinations within this circuit—see Fig. 14.




FIG. 14.

In this manner the flat three-resistor circuit can be hierarchically decomposed in at least three different ways: (1) as three separate resistors; (2) as the series combination of the upper two, in series with the lowest; (3) as the uppermost, in series with the series combination of the lower two. In a given circumstance one point of view may be more useful than another. Moreover, the various points of view do not fit neatly into a single hierarchy. They can be organized into an almost-hierarchy.

A complex example

The preceding examples show how alternative points of view can help us avoid algebra in simple cases. Here we show a more substantial problem whose solution can be facilitated with redundant descriptions. Consider a problem which we have used in class.¹⁰

¹⁰ Sussman has been in charge of the introductory subject in Electrical Engineering at MIT, Introductory Network Theory. He finds the relationship between teaching introductory engineering subjects and research in Artificial Intelligence rewarding. Teaching novel ideas to humans is like teaching them to machines. Often a new way of explaining an old idea becomes a program and the reverse is also true. Observing expert circuit analyzers solve networks, so that students could be taught their expert methods, led to the development of analysis by propagation of constraints. This became EL. EL style analysis is now being taught to our students.

WEEI-AM (590 kHz) needs a 'dummy load' for use in the maintenance of their 50 kW transmitter. A dummy load is a resistor, connected to the output of a transmitter in lieu of an antenna, to allow adjustments and measurements to be made on the transmitter under controlled conditions without radiating the test signals (which might interfere with other services). The only resistor they have which is capable of dissipating 50 kW is a 1 Ohm graphite block in a drum of oil (used to carry off the heat). The transmitter, however, needs a 50 ohm resistive load—it is to be adjusted to operate into a 50 Ohm antenna. Design a circuit to match the transmitter to the 1 Ohm resistor.
(Hint: You need just one inductor and one capacitor.)

The first difficulty here is understanding the problem. Ignoring the difficulties of natural language and the side issues which make the problem 'real' to the student, how can we formally specify the desired construct?

Just as we have summarized the behavior of the antenna at 590 kHz as a pure resistance of 50 Ohms we specify the behavior of the desired network in terms of an equivalent circuit. We want to construct a 2-port network, N , such that if it has a 1 Ohm resistor connected to its second port, the resulting 1-port network, N' , is equivalent to a 50 Ohm resistor at 590 kHz. A further reasonable constraint is that N dissipates no power itself.





FIG. 15. $R_1 = 50$, $R_2 = 1$

The failure of analysis

Let's ignore, for the nonce, the problem of finding a candidate topology. Suppose that somehow we are told that an appropriate circuit for N is an L-network:



The problem is now reduced to one of determining values of L and C for which N displays the behavior desired.

Since the frequency is a constant of the problem, we can simplify things by working with the reactances rather than with frequency-dependent impedances. We define:


$$X_L = L\omega, \quad X_C = \frac{1}{C\omega}$$

Thus the real problem is one of finding the reactances.

We analyze the network N' to determine the impedance it presents.

$$R_2 = 1\Omega$$

$$\omega = 2\pi f \simeq 3.7 \times 10^6 \frac{\text{rad}}{\text{sec}}$$



We can combine the impedances in series and parallel to get:

$$Z = \frac{-X_C j(X_L j + R_2)}{R_2 + (X_L - X_C)j}$$

Simplifying, we obtain:

$$Z = \frac{X_C X_L - R_2 X_C j}{R_2 + (X_L - X_C)j}$$

Separating the real and imaginary parts, we get:

$$Z = \frac{R_2 X_C^2}{R_2^2 + (X_L - X_C)^2} + \frac{-X_L^2 X_C + X_C^2 X_L - R_2^2 X_C j}{R_2^2 + (X_L - X_C)^2}$$

Our problem statement required the impedance to be R_1 . Using our given numbers for R_1 and R_2 , we get:

$$50 + 0j = \frac{X_C^2}{1 + (X_L - X_C)^2} + \frac{-X_L^2 X_C + X_C^2 X_L - X_C j}{1 + (X_L - X_C)^2}$$

Our complex equation reduces to two nonlinear equations in the two unknowns (X_L and X_C). Most engineering students can solve these, but we have only two undetermined parts. Suppose we were trying to design a T-section filter? The problem is hopeless, even with a powerful algebraic manipulator,¹¹ without some better idea. Analysis for voltages and currents is often easy because of linearity but the component values are nonlinearly related to the specifications.


Knowing how it works

The solution to this problem lies in the knowledge of how an L-network accomplishes the goal of matching the impedances. Each part in the L-network contributes to accomplishing the goal. The key idea can be seen by looking at the

¹¹ Although there are some very powerful algebraic manipulators available [21], it is unreasonable to expect them to help us out by symbolically inverting the equations of analysis. Suppose that the problem was an 8-section lattice filter (32 part values to determine—nonlinearly related to each other and to the specifications)? In general, it is more helpful to have knowledge which tells us how to avoid manipulation than it is to have powerful manipulators.

Believing in the ultimate power of mathematical manipulation is one of the most common difficulties encountered by students learning electrical circuit analysis. Students often grind out 'impossible' algebra in the course of solving a homework problem, even though a little thought will reveal an algebraically feasible approach which depends upon a small insight into the operation of the network being analyzed. They then complain that we give them too much homework!


parallel equivalent of a series combination. Consider just the inductor-load combination:



If we look at the admittance rather than the impedance, we obtain:


$$Y'' = \frac{1}{Z''} = \frac{1}{R_2 + X_L j} = \frac{R_2}{R_2^2 + X_L^2} - \frac{X_L}{R_2^2 + X_L^2} j$$

This admittance is the sum of two terms. We can view this as a parallel combination of two admittances: a frequency-dependent conductance and a frequency-dependent inductive susceptance. At any particular frequency there is a resistor and an inductor whose parallel combination is equivalent to N'' :



Note that the equivalent parallel resistance R'' can be made larger than R_2 by adding in inductive reactance. In fact, we can make this apparent parallel resistance the 50 Ohms (R_1) required by the problem.

Now we see the *purpose* of the parallel capacitor in the L-network. If we use it to resonate out the parallel inductance L'' at the specified frequency, all that will be left is the 50 Ohm parallel resistance we constructed with the series inductor and load resistor combination.



(The parallel combination of the capacitor and the apparent inductor is equivalent to an open circuit at the resonant frequency!)

With this understanding, the problem of finding appropriate values of L and C is significantly simplified. We can have algebraic manipulators, both human and mechanical, which can solve the problem.

If we set $R'' = R_1$ we can easily solve for X_L in terms of R_1 and R_2 at our desired frequency.

$$X_L = \sqrt{R_1 R_2 - R_2^2}$$

We then plug this result in to get the net parallel inductive reactance of network N'' .

$$X_{L''} = \frac{R_1 R_2}{\sqrt{R_1 R_2 - R_2^2}}$$

We must provide an equal capacitive reactance to balance this inductive reactance at the desired frequency. This condition determines the value of C .

$$X_C = \frac{R_1 R_2}{\sqrt{R_1 R_2 - R_2^2}}$$

Notice how smoothly the equations solve themselves. The trick is not being good at algebra but rather knowing what algebra to do. Each component in a circuit is used to provide for only part of the goal of that circuit. It is necessary to look at each part in its restricted context. In electrical circuits, this is often conveniently specified by means of equivalent circuits.

Power and phasors

Another point of view which explicates the operation of our L-network involves power arguments and reasoning with phasors. We intend that N be such that N'' looks like a 50 Ohm resistor R_1 . But no power is dissipated in N ; in fact all of the power apparently dissipated in R_1 is really dissipated in the 1 Ohm resistor R_2 .

Now suppose the RMS voltage across N' is $|V_1|$. Then the power dissipated is $|V_1|^2/R_1$. But this power is really going into R_2 hence the RMS voltage across R_2 is:

$$\sqrt{\frac{R_2}{R_1}} |V_1|$$

R_2 and L are in series so the RMS voltage across L is:

$$\sqrt{|V_1|^2 - \frac{R_2}{R_1} |V_1|^2}$$

In addition, the current through R_2 (and L) is:

$$\sqrt{\frac{1}{R_2 R_1}} |V_1|$$

Hence, we can derive the reactance:

$$X_L = \sqrt{R_1 R_2 \left(1 - \frac{R_2}{R_1}\right)} = \sqrt{R_1 R_2 - R_2^2}$$

We could continue to extract the capacitive reactance . . . but later.

Expressing this with slices

The moral of this story is that the activity of assigning component values for the parts of a circuit whose topology is known is not a simple task. To be able to

accomplish it without unobtainable algebraic power requires an understanding of how the circuit works. This combines a diverse set of reasoning strategies and overlapping points of view. Some of these points of view are expressed by useful electrical port equivalences, and some are more global statements about the identity of parameters. We now use these techniques to present the complete set of slices which constitute the rationale for the design of an L-network impedance matcher.




FIG. 16. L-network and explanatory slices.

We have tried to capture in Fig. 16 all of the basic relationships which explain how the L-network is intended to work. We see the following facts indicated:

The network is expected to be connected on its right to a circuit which is equivalent to a resistance R_2 .

The network presents a resistance R_1 to the port on its left.

The series combination of L and R_2 is equivalent to the parallel combination of L'' and R'' .

R'' has the same value as R_1 .

The parallel combination of L'' and C is equivalent to a current source of zero current (an open circuit).

The power dissipated in R_1 is the power dissipated in R_2 .

Some of the slices express information not available from the basic circuit diagram. For example, the slices a and b , which attach the equivalent resistances R_1 and R_2 to the circuit form the specifications of the L-network. The power relationship and the parallel equivalence are redundant as they could be deduced by a sufficiently powerful circuit analyzer.

Synthesis by analysis

Aided by the slices we can now determine the component values of the L-network by propagation. We terminate the L-network with a 1 Ohm resistor on its right and we specify that its input impedance be a 50 Ohm resistance. (See Fig. 17.)




FIG. 17. L-network and specifications.

Suppose we knew the voltage-amplitude, V_2 (relative to the ground indicated) at the top of the 1 Ohm load, then the current (amplitude) flowing down through the load is V_2 . This same voltage and current appear on the terminals of R_2 so we can deduce in one step that $R_2 = 1$ Ohm. This current also flows through the inductor, L ; it enters the left side of L . (Thus the parallel slice has this same current entering it.) Now the power dissipated in R_2 , $P_2 = |V_2|^2$. This power must be dissipated in R_1 . Suppose we also knew the voltage V_1 at the top of R_1 . Thus by a similar argument to the one above, we can deduce that $R_1 = 50$ Ohms. Since we know the power dissipated in this 50 Ohm resistance, we can deduce that the magnitude of the voltage at the top of R_1 is $\text{SQRT}(50) * |V_2|$.

At this point we know the voltage amplitude at the right of the inductor L , the current amplitude through it, and the magnitude of the voltage on its left. From this we deduce its reactance:

$$X_L = 7 \text{ Ohms.}$$

Using this we can get the full voltage amplitude on the left of L in terms of that on the right:

$$V_1 = V_2 * (1 + 7*j)$$

This voltage also appears at the top of R_1 . We use this to get the current through R_1 (which is the same as the current into N'). It is:

$$V_2 * (1/50 + (7/50)*j)$$

From this we compute the current into the capacitor:

$$V_2 * (-49/50 + (7/50)*j)$$

But we know the voltage across the capacitor and the current through it; hence the reactance

$$X_C = 50/7 \text{ Ohms.}$$

Note that we didn't use the parallel slice. Other interesting derivations pass through that route. The slices are redundant descriptions—the same truths from a variety of viewpoints. The way that the slices do their job, however, is *by providing redundant paths* for information to travel in the process of analysis. Thus, in the power argument above, the power dissipation slice, though redundant (It is derivable using only the local knowledge that inductors and capacitors dissipate no net power.), made it possible for enough information to get from the right side of the inductor to its left side to determine the inductive reactance.

Knowing the form of the answer

We have shown how slices can be used to express some of the knowledge of how a circuit is intended to work by describing alternate views of sections of the circuit. We have shown how this knowledge can help in the process of synthesis *if we know the topology and the correct slices to use*. Where does this information come from?

To answer this question in detail requires describing a theory of the design process. We can outline such a theory in general terms. The key idea is that of *knowing the form of the answer*. When presented with a simple problem, like a radio frequency impedance match, an engineer is likely to be able to retrieve such plausible solutions as the L-network and the transformer directly. If a more complex problem is posed, such as a 30 dB gain broadband amplifier with an input and output impedance of 600 Ohms, capable of output signals with 1 Volt of swing, the engineer does not retrieve a detailed answer but rather a high-level plan such as in Fig. 18.





FIG. 18. 2-stage plan.


This has the effect of decomposing the problem into a set of relatively independent subproblems. In fact, there are always complex interactions among the

parts which lead to bugs and debugging.¹² But the form of the answer is set. The problem is reduced to one of finding a feasible solution for each subproblem represented by the boxes. Each subproblem is constrained by the laws of electricity and the slices present on the plan. Each subproblem is further expanded until one gets to atomic problems.¹³

We believe that the slices are an integral part of the forms of solution for the different problem classes stored in the engineer's bag-of-tricks. When solutions to subproblems are combined, as above, the slices on the subsolutions are combined with the slices on the higher level plan. In addition, new patterns are formed by debugging of designs and by certain window optimizations. For example, if the solution of each of two subproblems leads to an inductor in series with a port,



and the two subsolutions are combined according to a Cascade plan it is recognized as an optimizable pattern and replaced by one inductor which performs the duties of both subinductors:



The old circuit remains as a set of slices explaining the new part!

Knowing the form of the answer is not unique to engineering problem solving.

¹² We believe that many bugs are just manifestations of powerful strategies of creative thinking—that creation and removal of bugs are *necessary* steps in the normal process of solving a complex problem. Following the work of Polya [26], recent research [12, 17, 33] predicated on this belief has resulted in the development of a paradigm for problem solving which we call *Problem Solving by Debugging Almost-Right Plans* (PSBDARP). We believe that the PSBDARP theory is a good foundation for building expert problem-solving systems for such diverse kinds of engineering as circuit design [35, 36] and computer programming.

¹³ The idea of successive refinement of plans appears as a key dogma of 'Structured Programming' [5, 9, 43], although it also appears in the Artificial Intelligence problem-solving literature. The idea of relaxation of a hierarchy of constraints comes from Freeman and Newell [15]. There are also versions of GPS [11] which were purported to do reasoning in a hierarchy of abstraction spaces. ABSTRIPS [28] showed how refinement of abstract plans could be used to guide a problem solver past problems which would otherwise be combinatorially explosive. Recently the NOAH system [29] has developed this idea to great depth.

It is very common in mathematics.¹⁴ For example, if we have a Homogeneous Linear Differential Equation with Initial Conditions we know that the form of the answer is $Ae \cos(wt + p)$. We can plug this in and solve for A, k, w, p . The algebra is, however, pretty rough. If we know how the parts of the answer are related to the problem—that w and k are functions of the differential equation only and A, p are derived from the effects of the initial conditions—things are algebraically more feasible. These facts can be represented as 'algebraic slices.'

4. Constraints and Algebra

It might seem that in writing down and using constraints we are just writing algebraic expressions in a different and not even convenient form. This is true to some extent, but it is not the entire picture. The constraint language is a kind of 'structured algebra' in that we have almost-hierarchical structures which allow us to manufacture and combine systems of algebraic constraints to make larger systems. Additionally, we will see that networks of constraints are more fundamental than the algebraic notation used to describe them, in that there may be many equivalent sets of equations which describe any given constraint diagram, but there is precisely one constraint diagram for any given set of equations.

Algebra has two components of interest. One is a *notation* for expressing constraints, and the other is a *set of transformation rules* which permit the derivation of consequences of the constraints.

The reader may have noticed that the constraint language is rather verbose. In defining the numerical constraints in the RESISTOR constraint definition, we had to write out eight lines:

```
(== (>> v t1) (>> sum av))
(== (>> v t2) (>> a1 av))
(== (>> a2 av) (>> product m))
(== (>> i t1) (>> m1 m))
(== resistance (>> m2 m))
(== (>> i t1) (>> a1 ai))
(== (>> i t2) (>> a2 ai))
(constant (>> sum ai) 0.0))
```

Using an algebraic notation, we might have written simply:

```
(>> v t1) - (>> v t2) = (>> i t1) * resistance
(>> i t1) + (>> i t2) = 0.0
```

Algebraic convention normally dictates that simpler variable names than $(>> v t1)$ be used. We have used the same long names in the algebraic equations not only

¹⁴ Perhaps the best example of a program which uses The Method of Knowing the Form of the Answer is SIN [25]. This symbolic integration program classified an expression to be integrated according to certain criteria which allowed it to select a general form for the answer and then solve for the details of the undetermined coefficients. SIN did not, however, use slices to help in the solution process.

for the sake of a 'fair' comparison, but because names of this form are critical to the almost-hierarchical composition of the constraints.

Algebraic notation is expression-oriented. It achieves conciseness through the use of functional composition. The relationship of arguments to function is implicitly expressed by position, permitting the elimination of explicit description of the connections. For example, in expressing an addition constraint in the constraint language, we explicitly mention the three 'pins' A1, A2 and SUM of the adder, and explicitly specify what they are linked to. In algebraic notation we simply write ' $a+b$.' The $+$ denotes an adder constraint. The three connections to the adder constraint are expressed by position. Whatever is to the left of the $+$ is connected to the A1 pin; whatever is to the right of the $+$ is connected to the A2 pin; and the SUM pin is connected to the pin of whatever operator of which the entire expression stands as an argument. If we write ' $(a+b) * c$,' then the SUM pin of the adder constraint is connected to the M1 pin of the multiplier constraint.

The connections indicated in an algebraic expression form a tree, with connections to identifiers at the leaves, and a single loose connection coming 'out the top.' To connect the loose ends of two expressions, we use the equality symbol $=$. An equation, unlike an expression, is a constraint network with no loose ends (or rather, the loose ends are explicitly indicated by the presence of identifiers).

The advantages of algebraic notation are obvious. What do we pay for conciseness? Asymmetry. To a mathematician, the following three statements say the same thing:

$$E = K + U \quad K = E - U \quad U = E - K$$

Each 'binary' operator must in principle have three forms, one for each of the three connections that may come 'out the top' as opposed to being connected to 'arguments.' (In the case of addition and subtraction two forms suffice because one of them is commutative. Consider, however, the triplet of equations

$$x = y^z \quad y = \sqrt[z]{x} \quad z = \log_y x$$

each of which expresses the same constraint in a different form. More generally, an n -ary operator with no symmetries on its pins requires $n+1$ distinct forms.) The constraint language preserves symmetry of expression and economizes on primitive notations: both additions and subtractions are expressed by the same constraint. The cost of this symmetry and economy in the constraint language is verbosity.

Algebra provides rules for transforming among equivalent representations of a constraint. Given this, in some sense it doesn't matter how they are written down. Equations express truth.

Given a set of equations, we can easily draw the (unique) constraint diagram which they represent. For example, for the set

$$2 * X + 3 * Y = 3 \quad 5 * X + 4 * Y = 7$$

we draw the diagram of Fig. 19.




FIG. 19.

Each equation represents a piece of the network which is an unrooted tree, the result of connecting two rooted trees at their roots. The *purpose of the identifiers* X and Y in the equations is to *cut the loops in the network* in order to make them printable as (necessarily tree-like) expressions. When the separate tree-like pieces are printed, the identifiers indicate the cross-branch and inter-tree connections.

For any given network, there may be more than one way to cut the loops. For example, we could use another variable U —see the diagram of Fig. 20.




FIG. 20.

We can then write this cut of the network as

$$6 = U + 3 * ((7 - 5 * (U/2))/4)$$

Even for a given assignment of identifiers, one can write many different equations by choosing which (two-ended!) connection to use for the equality; the two parts of the equation tree are considered to be rooted at the ends of the chosen connection. Using the connections marked [a]–[e] we get these equations:

$$[a] \quad U = 2 * ((7 - 4 * ((6 - U)/3))/5)$$

$$[b] \quad U/2 = ((7 - 4 * ((6 - U)/3))/5)$$

$$[c] \quad 5 * (U/2) = (7 - 4 * ((6 - U)/3))$$

$$[d] \quad (7 - 5 * (U/2)) = 4 * ((6 - U)/3)$$

$$[e] \quad 3 = (6 - U)/((7 - 5 * (U/2))/4)$$

All of these equations represent the *same* network of relationships.

It is of course possible to use more identifiers than necessary. One (U) suffices for our example, and it could have been put in any of six places. We originally derived the diagram from two equations in two identifiers X and Y. We could use more than two if we desired to, for some reason; we could also assign more than one identifier to a connection, and equate the redundant identifiers. Taken to the extreme, this approaches the connection specifications of the constraint language, wherein each quantity has multiple names which are equated.

In some diagrams a single connection may have more than two 'ends'; i.e. the same quantity is constrained in more than two ways. This occurs in the RESISTOR definition above. Now the implicit connections in expression notation have only two ends. When a diagram with many-ended connectors is to be printed as a set of equations, such connectors *must* have identifiers assigned to them, because only identifiers can express multiple-ended connections in algebraic notation. After such identifiers have been assigned, additional ones may or may not be needed to cut remaining loops.

The various ways of algebraically notating a set of relationships are equivalent. A network of relationships simply exists. It is in this sense that we mean that constraint diagrams are more fundamental than sets of equations. Many algebraic laws (but not all) simply provide for transformations among sets of equations which represent the same diagram.

When we begin to use the network for some computational purpose, however, then the point of view becomes important. If we wish to compute some 'output' value(s) given certain 'input' values, where all the values are related by some network of constraints, we must organize an information flow within the network. Intermediate computations must proceed along an acyclic path from inputs to output(s).

So-called 'algebraic' programming languages find algebraic *notation* a convenience because the tree implicit in the functional composition is guaranteed acyclic, and there is an asymmetry to each connection (from 'out the top' to 'into an argument slot') which can be construed to imply the direction of information flow. The notation therefore requires the programmer not only to express relevant relationships but also simultaneously to specify the particular computational use to which the relationships will be put. Thus, for example, Ohm's Law is the equation " $V = IR$ ", but the user of an 'algebraic' programming language cannot simply write this algebraic equation; he must put it in a particular form (for example, $I := V/R$) which explicitly directs the computation.

This may serve as a definition of the distinction between 'imperative' and 'declarative' languages. An imperative language requires the programmer to explicitly break loops in the network of constraints, and to explicitly organize the information flow within the network. A declarative language requires only the statement of relevant relationships, and the computational organization is specified separately or performed more or less automatically.

Besides requiring the programmer to specify the computational flow as a tree,

most 'algebraic' programming languages require expressions to denote single results.¹⁵ For example, there is generally no way to define an integer division "function" which returns both quotient and remainder (which is a shame, because typical algorithms for computing either actually compute both, particularly those used in many contemporary computers) and embed a call to such a "function" in an expression. This is an inherent property of trees.

We pointed out earlier that algebra provides transformation rules for "shifting perspective" on a relationship, for example transforming $E = K + U$ into $K = E - U$. Not all algebraic transformations are this simple. Some do not merely view the same network in a different way, but produce a different network which has the same meaning by virtue of non-trivial relationships between the meanings of the component constraints. A good example of this is the distributive law of multiplication over addition:

$$X = AB + AC \Leftrightarrow X = A(B + C)$$

The networks for these two equations look as shown in Fig. 21. These networks




FIG. 21.

are topologically distinct: one has a loop, and the other does not. Consider now trying to compute A given X , B and C . The second network can be used straightforwardly, because a spanning tree can be imposed on it, *rooted at A*, with *all leaves given*, specifying a complete computational flow. Such a tree cannot be imposed on the first network. Any spanning tree rooted at A has A as a leaf (a spanning tree rooted at a connection extends out only one end of the connection), and so is unsuitable for computing A .

This is why the voltage divider failed to compute the midpoint voltage given the endpoint voltages and the two resistances. The constraints imposed by the

¹⁵ There are a few languages which permit a routine to return more than one value (other than by side effect of a reference parameter as in FORTRAN). In SL5 every expression implicitly carries two values: an 'ordinary' value and a success/failure flag. Languages such as POP2 and FORTH are stack-based, and routines both receive arguments and return values on the stack. APL and LISP do *not* have multiple-valued expressions. A common trick is to construct a data structure (an array or a list) which contains several distinct values to be returned; but the called routine must explicitly construct this structure, and the caller must then explicitly decompose it.

resistor definition contained loops which prevented the system from finding an effective computational flow within the network. Imposing the SERIES-RESISTORS equivalence introduced extra paths which provided an effective spanning tree for the network in terms of the particular givens. (See Fig. 22.)

Similarly, when we 'solve an equation' by using the distributive law, we in effect impose another piece of network on the original network which preserves the semantics of the relationships but permits a spanning tree to be found.

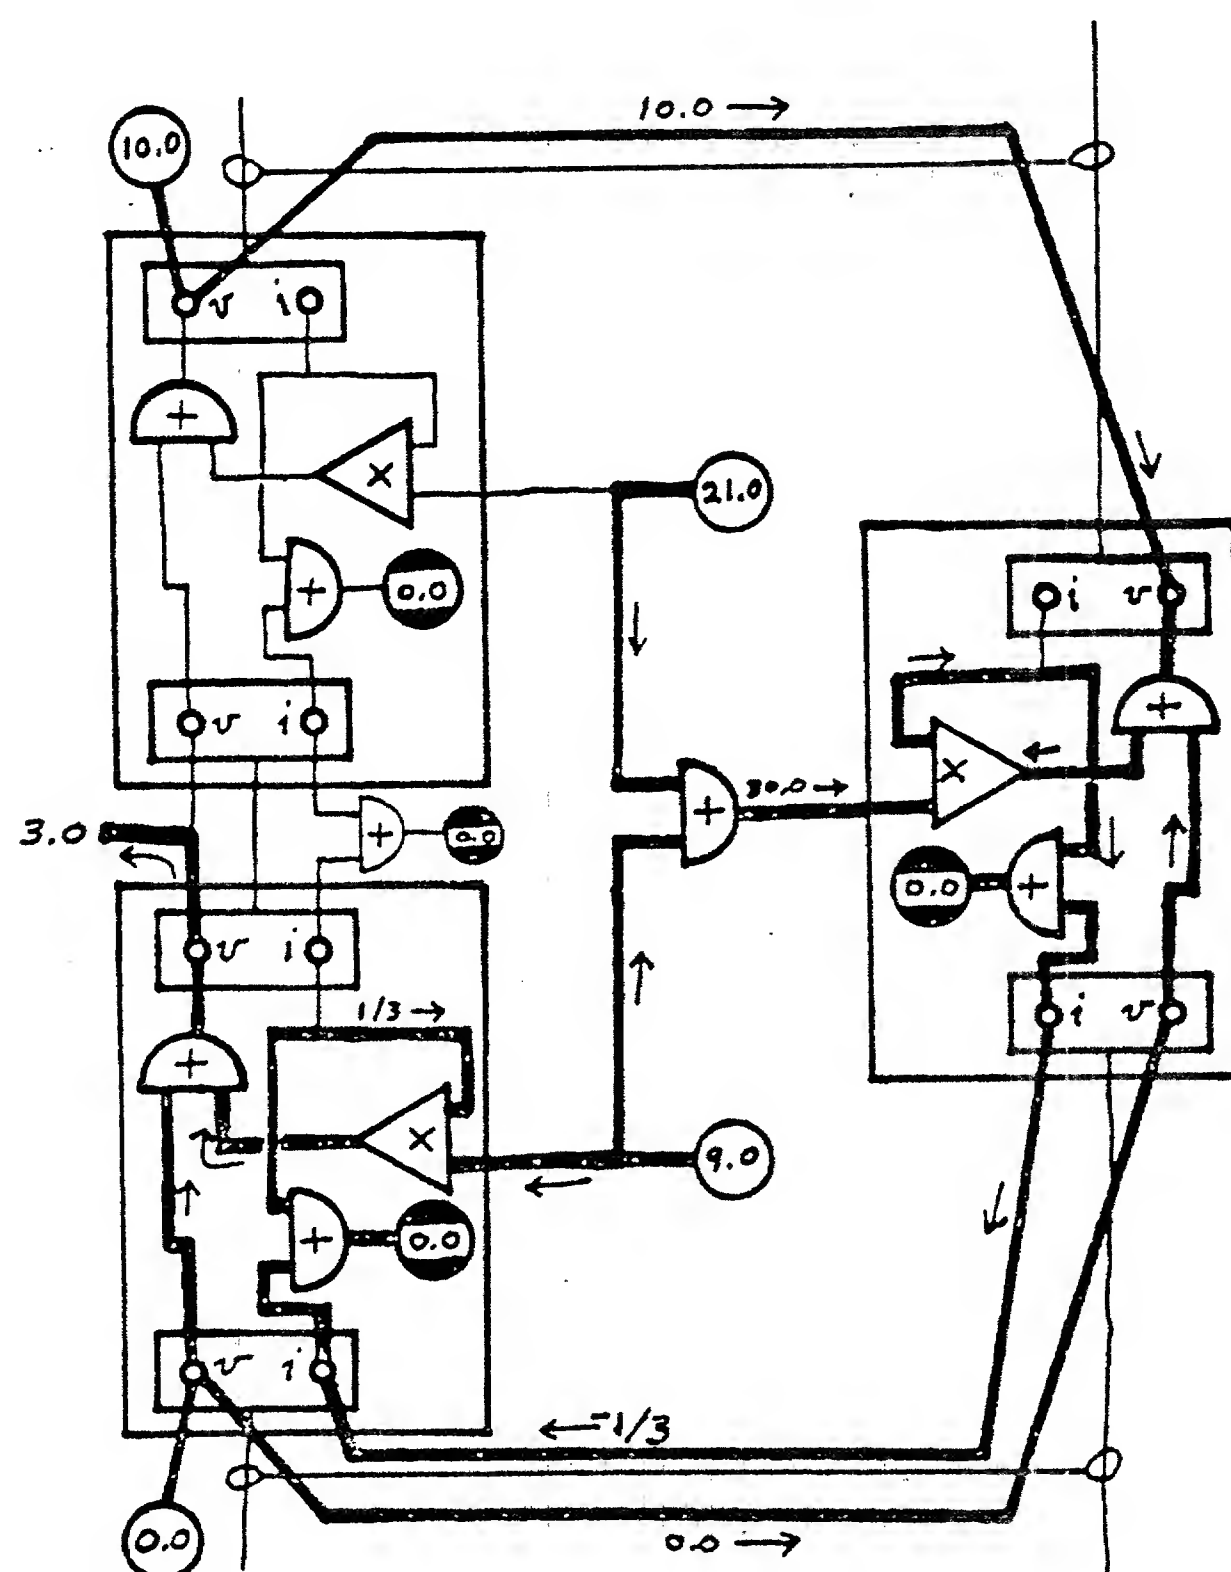


FIG. 22. Constraints for series resistors.

5. Computing with Constraints

All of the computations we have performed with constraint networks can be described as simple propagation of known values. Values are propagated through primitive constraints and through equalities by 'one-step deductions.' Primitive constraints are processes containing local cells which 'continually' monitor the cells. A cell may at any time have a value associated with it. A primitive constraint will notice if any of its cells gets a value, and if it can locally determine the value of one of its cells from the values of its other cells, it will associate the deduced

value with that cell. A cell can also get a value by being equal to another cell which gets a value. Finally, a cell can be arbitrarily assigned a value by the user or by a **CONSTANT** declaration (cf. **RESISTOR** above).

If a cell has an associated value, it also has associated with it a reason.¹⁶ If the value was obtained from other cells through a primitive constraint, then the reason mentions those other cells and the particular rule of the primitive constraint used to make the deduction. If the value was obtained through an equality with another cell, the reason mentions that cell. If the value was arbitrarily assigned, the reason points to the user or to the system (**CONSTANT**). These reasons ('dependencies') are easily constructed as the deductions are made; they are essentially simple markers indicating the direction of information flow within the network. These markers can be traced at any time for such purposes as finding the antecedent assumptions of a deduced value, or locating the consequences of a value. These may be in turn used for explanation or incremental forgetting of a user-supplied assumption and all of its consequences (as we have seen). In addition, if one traces from a value back to its antecedent assumptions, the part of the network thus traced, with the directionality of path imposed by the dependency markers, constitutes an algebraic formula for computing the value.¹⁷

It is possible that a cell be assigned a value for more than one reason. If two values collide at a cell they must be the same. If they are not, and if the network represents a satisfiable set of constraints, there is an inconsistent set of assumed values. The dependencies can be used to determine which subset of the set of assumed values is inconsistent. This can be helpful in dealing with very large networks with


¹⁶ TOPLE [23] was an early attempt to record the interactions among deductions for the purpose of maintaining consistency in a data base when newly introduced facts conflicted with existing ones. The SRI Computer Based Consultant [13] made use of dependencies to determine the logical support of facts, but did not use them to control search. MYCIN [6, 31, 32] used dependency information to produce explanations, but did not use it for any control purposes. EL [40] used dependencies to produce explanations and also to limit the recomputation required in response to incremental changes in the assumptions. Stallman and Sussman described a general means of limiting combinatorial search by the analysis of dependency chains and the use of multiple logical supports for facts [38]. This method, called 'dependency-directed backtracking,' was independently discovered by Jim Stansfield (unpublished Ph.D. thesis draft, University of Edinburgh—the idea was edited out of the final thesis for unknown reasons). This method is superior to the more familiar 'chronological backtracking,' introduced by Floyd [14] and best known for its use in AI languages beginning with **PLANNER** [18] and **Micro-PLANNER** [41], in that it avoids the irrelevant dependencies assumed by the system on the basis of accidental chronological orderings of assumptions [38, 39]. Doyle [10] developed a portable 'Truth Maintenance System' which encapsulates a careful theory of dependencies, dependency-directed backtracking and non-monotonic inference. A simpler system for dependencies and backtracking was developed in [22]. McDermott and Doyle have produced an elegant semantic theory of non-monotonic inference which is described in [24]. Phil London of University of Maryland has used dependencies to aid in keeping track of assumptions in the world model used by a problem-solving system [20].

¹⁷ The idea of extracting procedures for computing specific values from a constraint network is being pursued by Richard Brown [3]. Borning [1] has also done work in this area. (See also footnote 6, p. 11.)

many assumed values, where any one contradiction typically depends on only a small subset of the many assumptions.

In a satisfiable constraint network, the values which will be associated with cells are independent of the order of propagations through the network. Thus we may think of them as occurring in parallel, though they are implemented by means of a queue or agenda in our current system.

As we have seen, there are simple networks for which propagation fails to assign all of the values one might desire. For example:



We have already discussed how redundant descriptions may be used to bypass such loops by superimposing an equivalent tree structure. Another, related strategy involves the use of symbolic algebraic manipulation. If we enable the primitive constraints to propagate symbolic expressions as well as numerical values, we can use powerful algebraic manipulators to propagate past loops. If for example, in the problem above, we put the symbolic value X on (\gg A1 ZAP) then ZAP may deduce that its A2 is $5 - X$. 'Simultaneously' FOO may deduce that its M1 is $X/1$. These values are on cells declared to be equal; thus an equation may be formed, $X/1 = 5 - X$, and solved algebraically. Once the value of X has been determined, then the numerical values for the propagated symbolic expressions can be computed (or at least expressions involving X and possibly other variables can perhaps be simplified).

Algebraic techniques of this sort were used in such systems as EL, ARS, and SYN.¹⁸ It is easy to see, however, that propagation of symbolic expressions in effect make copies of tree-like portions of the network; each expression is a history of the portion of the network through which its pieces were derived. The algebraic techniques applied to these expressions might as well have been applied directly to the network. Instead of applying the distributive law to an algebraic expression, for example, one might as well just attach that extra point of view to the network. The implications of this idea have just barely begun to be explored.

¹⁸ We have used propagation of algebraic expressions in our work on analysis and synthesis of circuits. (See also footnote 2, p. 2.) Although we have used algebraic manipulation techniques (extracted from MACSYMA [21]) far more powerful than people generally use, we have found that these techniques are not, by themselves, powerful enough to solve many interesting problems which people can solve. People generally solve these problems by organizing the solution so that only simple algebra is required (often by using canned theorems) whose proofs perhaps required tremendously complex algebra (done once and for all by a clever person whose name is probably attached to the theorem), but whose application does not require complex algebra). To avoid tremendous symbolic computations, computers must also have good methods of avoiding doing all except the most trivial algebra problems. We have used multiple redundant descriptions [35, 36] to encapsulate ways of looking at a problem which are organized so as to obviate the need for extensive algebraic manipulation.

ACKNOWLEDGEMENTS

We would like to thank the many people who helped us. Original ideas about propagation and grey boxes were developed jointly with Richard M. Stallman. Significant ideas have also come from Drew McDermott, Allen Brown, Johan de Kleer, Jon Doyle, and Marvin Minsky. Danny Bobrow, Allen Borning, and Randy Davis made suggestions which improved the presentation. One of us (Steele) is supported by a Fannie and John Hertz Foundation graduate fellowship.

REFERENCES

1. Borning, A., ThingLab—an object-oriented system for building simulations using constraints, Proc. Fifth International Joint Conference on Artificial Intelligence (IJCAI-5), MIT, Cambridge (August 1977), 497–498.
2. Brown, A. L., Qualitative knowledge, causal reasoning, and the localization of failures, Ph.D. thesis, MIT, Cambridge (September 1975). Also MIT AI Lab Technical Report 362, Cambridge (March 1977).
3. Brown, R., Automatic synthesis of numerical computer programs, Unpublished Ph.D. thesis proposal, MIT, Cambridge (September 1978).
4. Church, A., *The Calculi of Lambda Conversion*, Annals of Mathematics Studies Vol. 6 (Princeton University Press, Princeton, 1941). (Reprinted Klaus Reprint, New York, 1965.)
5. Dahl, O. J., Dijkstra, E. W., and Hoare, C. A. R., *Structured Programming* (Academic Press, London, 1972).
6. Davis, R., Applications of meta level knowledge to the construction, maintenance and use of large knowledge bases, Stanford University, AI Lab Memo AIM-283 (Stanford, July 1976).
7. de Kleer, J., Local methods for localization of faults in electronic circuits. MIT AI Lab Memo 394 (Cambridge, November 1976).
8. de Kleer, J., and Sussman, G. J., Propagation of constraints applied to circuit synthesis, MIT AI Lab Memo 485, Cambridge (September 1978).
9. Dijkstra, E., Structured programming, in: Ruxton, J. N., and Randell, B., Eds., *Software Engineering Techniques* (Nato Scientific Affairs Division, Brussels, Belgium, 1970).
10. Doyle, J., Truth maintenance systems for problem solving, M.S. thesis, MIT, Cambridge (May 1977). Also MIT AI Lab Technical Report 419, Cambridge (January 1978).
11. Ernst, G. W., and Newell, A., *GPS: A Case Study in Generality and Problem-Solving* (Academic Press, New York, 1969).
12. Fahlman, S., A planning system for robot construction tasks, M.S. thesis. MIT, Cambridge (May 1973). Also MIT AI Lab Technical Report 283, Cambridge (May 1973).
13. Fikes, R. E., Deductive retrieval mechanisms for state description models, Stanford Research Institute, Menlo Park, California, AI Technical Note 106, (July 1975).
14. Floyd, R. W., Nondeterministic algorithms, *J. ACM* (October 1967).
15. Freeman, P., and Newell, A., A model for functional reasoning in design, *Proc. IJCAI II* (1977).
16. Goldberg, A., and Kay A., (Eds.), *Smalltalk-72 Instruction Manual* (Xerox Palo Alto Research Center, SSL 76-6, Palo Alto, California, 1976).
17. Goldstein, I., Summary of MYCROFT: a system for understanding simple picture programs, *Artificial Intelligence* 6 (1975), 249–288.
18. Hewitt, C. E., Description and theoretical analysis (using schemata) of PLANNER: a language for proving theorems and manipulating models in a robot, Ph.D. thesis. MIT, Cambridge (January 1971). Also MIT AI Lab Technical Report 258, Cambridge (April 1972).
19. Lacombe, J-C., Artificial Intelligence in computer-aided design: the tropic system, Stanford Research Institute, AI Center, Menlo Park, CA, Technical Note 125 (1976).
20. London, P. E., Dependency networks as a representation for modelling in general problem solvers, Ph.D. thesis, University of Maryland, Dept. of Computer Science, College Park, MD., Technical Report 698 (September 1978).

21. Mathlab Group, *MACSYMA Reference Manual, Version Nine* (MIT Laboratory for Computer Science, Cambridge, December 1977).
22. McAllester, D. A., A three-valued truth maintenance system, MIT, AI Lab, Cambridge, Memo 473 (May 1978).
23. McDermott, D. V., Assimilation of new information by a natural language-understanding system, B.S./M.S. thesis. MIT, Cambridge (1973). Also MIT, AI Lab, Cambridge, Technical Report 291 (February 1974).
24. McDermott, D. V., and Doyle, J., *Non-Monotonic Logic I, Artificial Intelligence*, To appear.
25. Moses, J., Symbolic integration, MIT, Cambridge, Project MAC, Technical Report 47 (1969).
26. Polya, G., *Mathematical Discovery: On Understanding, Learning, and Teaching Problem Solving, Vols. I and II*. (Wiley, New York, 1972).
27. Rich, C., Shrobe, H. E., and Waters, R. C., Computer aided evolutionary design for software engineering, MIT, Cambridge, AI Lab, Memo 506 (January 1976).
28. Sacerdoti, E., Planning in a hierarchy of abstraction spaces, *Proc. IJCAI III* (August 1973).
29. Sacerdoti, E., A structure for plans and behavior (Elsevier, New York, 1977).
30. Simon, H., *The Sciences of the Artificial* (MIT Press, Cambridge, 1969).
31. Shortcliffe, E. H., MYCIN—A rule-based computer program for advising physicians regarding antimicrobial therapy selection, Stanford University, AI Lab, Memo AIM-251 (October 1974).
32. Shortcliffe, E. H., MYCIN: Computer-based medical consultations (Elsevier, New York, 1976).
33. Sussman, G. J., *A Computational Model of Skill Acquisition*. Ph.D. thesis. MIT (1973). Also *A Computer Model of Skill Acquisition* (Elsevier, New York, 1975).
34. Sussman, G. J., The virtuous nature of bugs, *Proc. Artificial Intelligence Simulation of Behavior Conference* (Sussex, U.K., 1974).
35. Sussman, G. J., SLICES: At the boundary between analysis and synthesis. MIT, AI Lab, Cambridge, Memo 433 (July 1977). Also in: IFIP W.G. 5.2, Working Conference on Artificial Intelligence and Pattern Recognition in Computer-Aided Design.
36. Sussman, G. J., *Electrical Design—A Problem for Artificial Intelligence Research*, MIT, AI Lab, Cambridge, Memo 425 (June 1977). Also *Proc. IJCAI V* (August 1977).
37. Sussman, G. J., Holloway, J., and Knight, T. F. Jr., Computer aided evolutionary design for digital integrated systems, MIT AI Lab, (Cambridge, Memo 526, 1979).
38. Stallman, R. M., and Sussman, G. J., Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis, *Artificial Intelligence* 9 (1977), 135–196.
39. Sussman, G. J., and McDermott, D. V., From PLANNER to CONNIVER—A genetic approach, *Proc. AFIPS 1972 FJCC* (AFIPS Press, Montvale, N.J., 1972), 1171–1179.
40. Sussman, G. J., and Stallman, R. M., Heuristic techniques in computer-aided circuit analysis. *IEEE Transactions on Circuits and Systems* CAS-22 (11) (November 1975).
41. Sussman, G. J., Winograd, T., and Charniak, E., *Micro-PLANNER Reference Manual*. MIT, AI Lab, Cambridge, AI Memo 203A (December 1971).
42. Sutherland, I. E., SKETCHPAD: A man-machine graphical communication system, MIT, Lincoln Laboratory, Cambridge, Technical Report 296 (January 1963).
43. Wirth, N., Program development by stepwise refinement. *C. ACM*, 14 (4) (April 1971).

Received 17 September 1979; revised version received 25 September 1979

